

SYSTEM AND METHOD FOR INCREASED NETWORK SECURITY

Field of the Invention

This invention relates to methods and systems for increasing security over the Internet. More particularly, this invention relates to a system and method for providing public access to websites while providing increased security for private databases and applications that are accessed by the website in order to respond to user inquiries.

Background of the Invention

The internet may be loosely described as a public network, or a collection of networks, consisting generally of a collection of distributed IP devices, *e.g.*, PCs, printers, servers, etc., each of which have a distinct IP address, *e.g.*, 1.1.1.1, 1.1.1.2, ... 1.1.1.10. *Fig. 1*. There also exist many private networks which themselves comprise collection of devices, each having a discrete IP address which theoretically may be identical to the IP address of a completely different device located on a different private network or located on a public network. Private networks are often connected to the public networks through a gateway or proxy). Devices on one private network can communicate or “address” a device on a public network through the gateway that acts as a proxy to the public network. If a device on a private network knows the IP address of device in the same private network or public network, that device can directly address the second device using the above principle.

Accordingly, in order to address a device on the internet, it is necessary to know its IP address. The websites that make up much of what is considered by the public to be “the internet” reside on web servers which have addresses on a public network (a public IP address). By contrast, most users access the internet through PCs residing on a private network. A gateway provides a gateway to the public networks (these gateways are commonly referred to as Internet gateways and carry two IP addresses, one on the private network and the other on a public network). While these devices on private networks are described above as having IP addresses, these IP addresses are usually private (they are not addressable by users outside of the private network on which it resides) and dynamic (the IP address for a particular device on a

private network is assigned automatically by a DHCP (dynamic host configuration protocol) server in order to reduce network management overheads created by conflicting IP addresses).

[0004] Both the browser in a user's PC and a website residing on a web server have port numbers for receiving communications into the correct application that runs in the device. However, in contrast to the private and dynamic port numbers assigned to browsers and other client applications in a user's PC, website port numbers are public and well known. The port number for websites on the http network is 80; the port number for websites on the https is 443; and the port number for websites on the FTP network is 21. Accordingly, when a user's PC browser sends a request to a website, it uses the website's IP address together with the appropriate port to address the website and accompanies the message with the PC's temporary dynamic IP address as well as the transient port number that corresponds to the browser which generated the request. The website, in responding to the request, directs the response to the dynamic IP address and port number that accompanied the request. See, e.g., Fig. 2.

[0005] The way that PCs with dynamic IP addresses on a private network can receive information from a website web server on a public network may be described as follows: the user types into a browser application running on his PC the public IP address for a website web server. The user's PC then directs a request to the website server at the specified address and well-known port number. The request generated by the user's PC includes its return address (its private dynamic address) and a "port number", which is also known as a "TCP socket," that will be open for a short specified time in order to receive the reply. The TCP socket identifies the particular application on the user's computer to which the reply should be directed -- any one user may have several browsers open on his PC, using each browser to communicate with different websites.

[0006] Although a device in the private network is able to address a device in the public network through a gateway (which serves as a proxy for the public network), a website cannot address the PC in the private network due to the fact that private IP is not a 'global address' and it is not addressable from the public network. To enable the website to send the reply back to the private PC, the gateway manages a temporary network address translation table (NAT) which gets built dynamically as communication takes place. When a request from private

network travels to the public network via the gateway, it forwards the request to the website with its own public IP address and a randomly picked port number as the return socket. Also, the gateway adds an entry to the NAT to map that socket to the return IP address and port number of the browser or other client application that generated the request. When the reply is received from the website, the gateway forwards it to the client application the generated the request after performing a lookup in the NAT. Once the communication is completed (TCP connection is closed), the related entry in the NAT is removed; therefore, no more data packets can travel into the private network. If someone attempts to initiate a TCP communication into the private network, the gateway does not allow it by virtue of the fact that there is no address mapping in the NAT (unless it has been specifically setup to act as a reverse proxy as described below).

[0007] In this fashion, a public website can always be addressed by a client application on a private network, while a private user's PC cannot be addressed from the Internet.

[0008] As discussed above, websites traditionally reside on "web servers" which have a static public IP address on a public network. Alternatively, the web server may be represented on the public network by a router or "reverse proxy" which directs inquiries to the web server which may be placed on a private network. See, e.g., Fig. 3. In this case, the proxy or router will map its public IP address to the private address of the web server, and redirect inquiries coming into the public IP address to the web server address. This method is also known as "reverse proxy" or "IP forwarding" or "protocol tunneling" with slight variation in implementation. In either case, the principle is to forward the request coming from a client application (or an Internet gateway of another private network) on the Internet to the web server.

[0009] Websites, in contrast, by virtue of the fact that they have public addresses, are subject to unauthorized access. As mentioned above, websites typically reside on a web server and comprise two primary functional units: a listening unit and a responder unit. The listening unit, which maintains an open line of communication with the public network and receives requests for information from other devices (users) located on the public network or on private networks with access to the public network. The responder unit contains the ASP pages, CGI applications, etc., in effect defining that information which is to be published or made available

for publication. When the responder receives a request for information, the responder typically accesses a memory, for example via a database application, containing public and sensitive information. Many websites also have access to private “source data” which may be used to generate the public or sensitive information for publication to authorized users. The responder unit serves as the gateway for determining which requesting devices are entitled to sensitive and/or public information in the website. The responder unit is typically designed so as not to give away, or “publish,” the private source data. Rather, it only uses the private source data to generate the public and/or sensitive information which is then published to authorized users via the website. In order to prevent unauthorized access to sensitive information and private source data that available to a website responder unit, network engineers design “firewalls” which will attempt to identify instances of unauthorized access to the website’s data sources. This is primarily done by blocking outside users or devices from initiating TCP/IP connections into the protected network through specific ports (known as “blocking incoming ports”). However, firewalls cannot fully close all incoming ports into the website because certain ports must remain open for the web server to function. Further, a firewall only blocks the initialization of a TCP connection (at the beginning of the TCP conversation) by inspecting traffic that targets a specific port. After the initialization, the traffic has to pass through the firewall in both directions with randomly assigned ports, and the firewall has to allow it to happen. Therefore, unlike in the case of a proxy, a firewall is unable to isolate the network from unauthorized incoming traffic. Each call comes from a ‘visitor’ of the Internet, which could potentially be a hacker.

Summary of the Invention

[0010] The present invention arises from the realization of the inventors that unauthorized users can circumvent firewall protection of a website’s data sources by downloading a software file onto the website web server which replaces or modifies the existing website responder unit. See Fig. 4. Such replacement responder unit would be different from the original responder unit at least to the extent that it would allow the unauthorized user unlimited access to the website’s data sources. In other words, the responder unit of the website, which controls access to the public and sensitive information and private source data contained in the website’s data sources, might be replaced by a responder unit of the unauthorized user’s

own design, which design could permit unrestricted access by the unauthorized user to all of the sensitive information and private source data to which the website has access.

[0011] The present invention relates to a system and method which prevents this and other types of unauthorized access to sensitive information and private source data is used to serve websites which are accessible to the public. According to the invention, the website web server located on a public network having a public IP address and known port number performs only the listening function. The responding function is located on a separate device on a private network with a private and dynamic IP address and having a randomly assigned port number. The responder has no listening sockets (open ports expecting to receive from client application) and therefore does not listen to the public network, and therefore is not accessible to unauthorized access, much in the way that a private user's PC is not accessible to unauthorized access. The web server having the listening function does not initiate connection with the device having the responding function because its private IP address is unreachable from the public network and unknown, even to the web server, and by virtue of the fact that there are no listening sockets to accept any requests. Instead, the communication link between the device having the responder function and the web server having the listening function is initiated by the device having the responder function. Much in the way that a private user's PC on a private network opens a line of communication with a web server on a public network by sending a transmission which includes the PC's dynamic IP address and port number, the responder establishes a single encrypted connection to the web server, with a private and dynamic IP address and port number for responsive communications. In this fashion, the device on the private network having the responder function has access to the sensitive information and private source data and uses that information to provide information to the listening function for publication, but neither the responder nor the data sources are accessible to unauthorized users because there is only a single connection between the responder and the listener which connection is initiated by the responder. Thus, separation of the listening and responsive functions into two separate devices, with the responder device located on a private network, protects the responder functionality from unauthorized access by virtue of the fact that all communication between the listener and responder is transmitted over the responder initiated single TCP/IP connection. Further, the communication is encrypted for enhanced security.

Brief Description of the Drawings

[0012] Figure 1 shows a schematic representation of a private network in communication with a public network.

[0013] Figure 2 illustrates the communication between the browsers on a private user's PC on a private network with web servers located on a public network.

[0014] Figure 3 is a schematic representation of a website on a private network represented on a public network by a router or "reverse" proxy.

[0015] Figure 4 is a schematic representation of a web server in which an unauthorized user has placed a substitute responder, giving that user access to sensitive information and private source data on the database

[0016] Figure 5 is a schematic representation of the invention showing a server having only the listening function of a website located on a public network and a separate device having the responder function located on a private network with a dynamic IP address and port number with a single encrypted connection between the responder device and a listening server which is established by the responder device.

[0017] Figure 6A is a schematic representation of a preferred embodiment of the invention.

[0018] Figure 6B is a schematic representation of another embodiment of the invention.

[0019] Figure 6C is a schematic representation of another embodiment of the invention.

Description of the Preferred Embodiment of the Invention

[0020] Referring to Figure 5, an IP device such as a server is located on a public network, with a public IP address and known port number. This device contains a listening function for a website which receives communications and inquiries for information from other devices located on the public network or on private networks in communication with the public

network. This listening function also transmit responses to the requesting devices. A second IP device is located on a private network and contains the responder function of a traditional website. This service has a dynamic private IP address. The responder function is in communication with a data source, e.g., a database. The data source may contain public and sensitive information and/or private source data necessary or useful in responding to inquiries received by the listening function. The responder function initiates a single connection with the listening function, transmitting its dynamic IP address and randomly assigned port number for responsive communications. This line of communication may be established in much the same way that a private PC on a private network opens a line of communication with a traditional website as described above in the background of the invention. This single line of communication is preferably encrypted. When the listening function receives a request for information from another device on the public network or on a private network in communication with the public network, the listening function “responds” to the opening communication established by the responder using the dynamic IP address and port number transmitted by the responder in the opening communication. The responder function processes the request, optionally accessing information in the data source(s), and communicates their response to the listening function, which in turn publishes the information to the requesting device using the requesting device’s IP address and port number. Preferably, the only connection between the device having the responder function and the public network is the single encrypted connection with the listening server, established by the responder device using its dynamic IP address and randomly assigned port number. Furthermore, the responder that initiates the TCP/IP connection with the listener intercepts all the incoming messages from the listener, and does not carry out any function that can be harmful to the system. In this fashion, there can be no unauthorized access to the device carrying the responding function, and hence no unauthorized access to the data base which may contain sensitive information and private source data.

[0021] Figure 6A shows a more preferred embodiment of the invention in which the responder unit is connected to a private network and has a private IP address and a randomly assigned port number. It can access other devices on the private network, but cannot be addressed by devices outside the private network. The listener unit is connected to a public network with a public IP address and known port number. The responder unit initiates

communication with the listener unit via an encrypted TCP/IP connection. Accordingly, neither the responder unit, other applications residing on the same device, or other devices to which the responder unit has access are susceptible to unauthorized access, either via the listener or via the responder's connection to its private network.

[0022] Figure 6B shows another embodiment of the invention in which the responder function is moved to another public address as opposed to a private network. In this embodiment, all the "incoming ports" are closed in the firewall to protect the responder from TCP/IP level attacks, and the responder device communicates with the listener device via an encrypted TCP/IP connection. Therefore, the responder unit is not susceptible to unauthorized access via the listener unit. However, the responder device is still addressable over the internet. Accordingly, the ultimate security of this embodiment depends heavily on the stability of the firewall.

[0023] Figure 6C shows yet another embodiment of the invention in which the connection from the responder unit is not addressable via any network, public or private, and in which its connection to the listening unit is non-IP based, for example, via IPX, ethernet (below IP level), serial communication, USB or Fireware.